Adoption of a SAML-XACML Profile for Authorization Interoperability across Grid Middleware in OSG and EGEE

G Garzoglio[*] [1], J Bester[2], K Chadwick[1], D Dykstra[1], D Groep[3], J Gu[4], T Hesselroth[1], O Koeroo[3], T Levshina[1], S Martin[2], M Salle[3], N Sharma[1], A Sim[4], S Timm[1], A Verstegen[3]

[1] Fermi National Accelerator Laboratory, Batavia, IL, USA
[2] Argonne National Laboratory, Argonne, IL, USA
[3] NIKHEF, Amsterdam, The Netherlands
[4] Lawrence Berkeley National Laboratory, Berkeley, CA, USA

Abstract. The Authorization Interoperability activity was initiated in 2006 to foster interoperability between middleware and authorization infrastructures deployed in the Open Science Grid (OSG) and the Enabling Grids for E-sciencE (EGEE) projects. This activity delivered a common authorization protocol and a set of libraries that implement that protocol. In addition, a set of the most common Grid gateways, or Policy Enforcement Points (Globus Toolkit v4 Gatekeeper, GridFTP, dCache, etc.) and site authorization services, or Policy Decision Points (LCAS/LCMAPS, SCAS, GUMS, etc.) have been integrated with these libraries.
At this time, various software providers, including the Globus Toolkit v5, BeStMan, and the Site AuthoriZation service (SAZ), are integrating the authorization interoperability protocol with their products. In addition, as more and more software supports the same protocol, the community is converging on LCMAPS as a common module for identity attribute parsing and authorization call-out. This paper presents this effort, discusses the status of adoption of the common protocol and projects the community work on authorization in the near future.

1. Introduction

The Open Science Grid (OSG) [1] and Enabling Grids for E-sciencE (EGEE) [2] are two large national Grid infrastructures operated respectively in the US and Europe. Both Grids federate dozens of computing centres that aggregate Peta-flops of computing resources and Peta-Bytes of storage capacity. Users are granted access to resources in virtue of their membership to those scientific communities that are partners with the Grid organizations. Because of the international nature of these communities, interoperability is considered high priority among Grid organizations.

In order to address interoperability of the EGEE and OSG authorization infrastructures, the two Grids have initiated a collaborative project with the Globus Toolkit [4] and Condor [5] groups. Mission of the authorization interoperability project [3] was to agree on a common protocol and implementation between resource-gateway middleware (Policy Enforcement Points – PEP) and the site authorization service (Policy Decision Point – PDP) (sec. 2). Accomplishing this mission, the project aimed at reaching three goals:

[*] To whom any correspondence should be addressed. E-mail: garzoglio@fnal.gov

1. share and reuse software developed for EGEE and OSG
2. give software providers (external to the Grid organizations) reference protocols to integrate with both Grids infrastructures
3. enable the seamless deployment of software developed by OSG or EGEE in the EGEE or OSG authorization infrastructures

In sec. 2 we introduce the Authorization Infrastructure models of OSG and EGEE. In sec. 3 we discuss how these goals have been met by the project and their limitations. We also discuss how the deliverables of the project drastically simplify the OSG infrastructure. In sec. 4 we talk about the deployment of the authorization interoperability infrastructure and the relative challenges, before concluding in sec 5.

2. Grid Authorization Infrastructure

The EGEE (now European Grid Initiative) and OSG security model is based on X509 end entity and proxy certificates for single sign-on and delegation. Both Grids trust a common set of Certificate Authorities (CA) to enable authentication across Grids. The set of common CA is agreed upon through joined organizations, such as the Joint Security Policy Group (JSPG) [14].

Access to resources is granted on the basis of the user identity and the user's membership to a community or Virtual Organization (VO). VOs are typically organized in hierarchical groups. In addition, members may have roles, such as *administrator*, within a certain group. The VO structure and user membership are organized and maintained through the Virtual Organization Membership Registration System (VOMRS) [15] and pushed to the Virtual Organization Membership Service (VOMS) [16] (step 1 & 2 in fig. 1). Alternatively, the VOMS administrative interfaces can be also used for this purpose. Depending on the Grid configuration, this information is also synchronized with the sites authorization services and augmented with appropriate identity mapping policies to allow privilege management at sites (step 3).

To access the Grid on behalf of a certain VO / VO group with a certain role, a user contacts VOMS to extend her certificate with a membership assertion signed by the server, the VO authority on member identity (step 4). This information is then pushed to the resource gateways (step 5) to access a variety of Grid resources (batch systems, storage systems, etc.).

In turn, a resource gateway, acting as a PEP, contacts the site authorization services, or PDP's, to grant the user access with the appropriate privileges (steps 7 to 10). The Authorization Interoperability profile [17] defines a common set of attributes to express and communicate these authorization assertions. It is based on the Security Assertion Markup Language (SAML) [7] profile of the eXtensible Access Control Markup Language (XACML) [8], both standards from the Organization for the Advancement of Structured Information Standards (OASIS). The profile extends XACML to standardized names, values, and semantics for common attributes and obligations. This defines a common vocabulary for OSG and EGEE to express the properties of the authorization assertion, such as the identity of the Subject requesting the authorization, the Resource targeted, and the Action requested within a given Environment. The details of the profile are described elsewhere [17].

3. Results of Authorization Interoperability

After more than one year of work, in 2008 the project released a profile document [6] and reference authorization call-out library implementations. EGEE implemented the authorization interoperability profile through the Site Central Authorization Service (SCAS) [9] PDP and corresponding PEP call-out modules; OSG implemented it through the Grid User Mapping Service (GUMS) [11] and Site Authorization Service (SAZ) [12] PDP's and Prima call-out module. With this infrastructure, shortly thereafter, the project demonstrated the interoperability of middleware developed by OSG in the EGEE authorization infrastructure and vice versa. This was a demonstration of goal 3 (sec. 1).

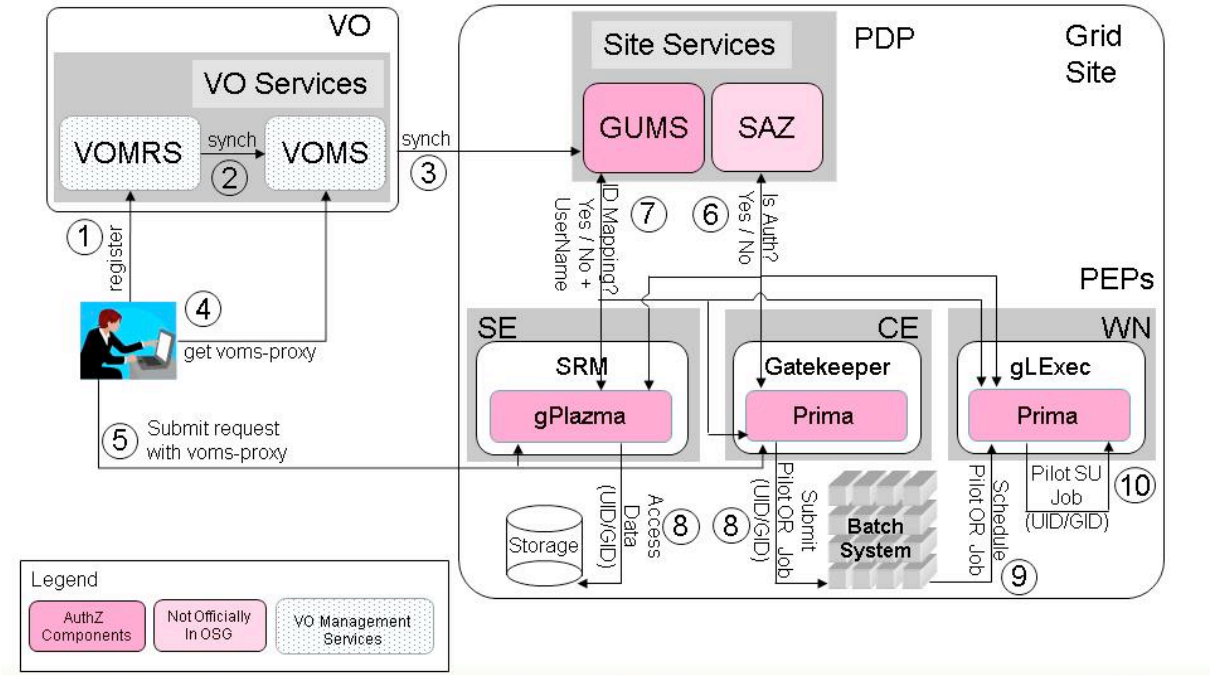

Figure 1. An architectural diagram of the Open Science Grid authorization infrastructure.

In 2009, the SCAS service was deprecated in favour of Argus, a new Grid authorization infrastructure. Argus based its authorization profile on an extension of the Authorization Interoperability profile. The main criticism to the non-extended profile was that some attributes, such as *subject-x509-id*, were not general enough according to the standards discussed in the Open Grid Forum (OGF). Attempts to collaborate on a common extension were not successful due to different timeline constraints. Because of the differences between the Argus and the Authorization Interoperability profile, interoperability between the two Grids is not enabled at this time.

While the interoperability goal is still in hold, in the same year, the profile documentation was used as a reference for the development of external projects. In particular, TechX Corp. developed the Scalable Virtual Organization Policy Management Environment (SVOPME) [13] using the profile to define authorization policies for Virtual Organization and Grid Sites. This experience is a good example of how goal 2 (sec. 1) for the project is fulfilled.

The common profile between EGEE (through SCAS) and OSG is naturally leading to the adoption of a common implementation for the authorization call-out modules. In 2010, OSG has started a program of work to adopt the EGEE framework for PEP authorization, the Local Centre Authorization Service and Local Credential Mapping Service (LCAS / LCMAPS) framework [10]. Coupled with the SCAS client plug in, the adoption of a single authorization framework has two main advantages:

1. it reduces the overall maintenance effort for our community, fulfilling goal 1 (sec. 1) for the project;
2. it drastically simplifies the OSG authorization infrastructure.

Figure 2 shows how the infrastructure is simplified using two architectural diagrams. In each diagram, PEPs, at the bottom, communicate with PDPs, at the top, through a set of authorization call-out frameworks and libraries (in the middle). Until 2010 (top diagram), because of the limited deployment of the authorization interoperability modules, PEPs communicate to individual PDPs using PDP-specific authorization modules. In addition, certain PEPs use PEP-specific authorization

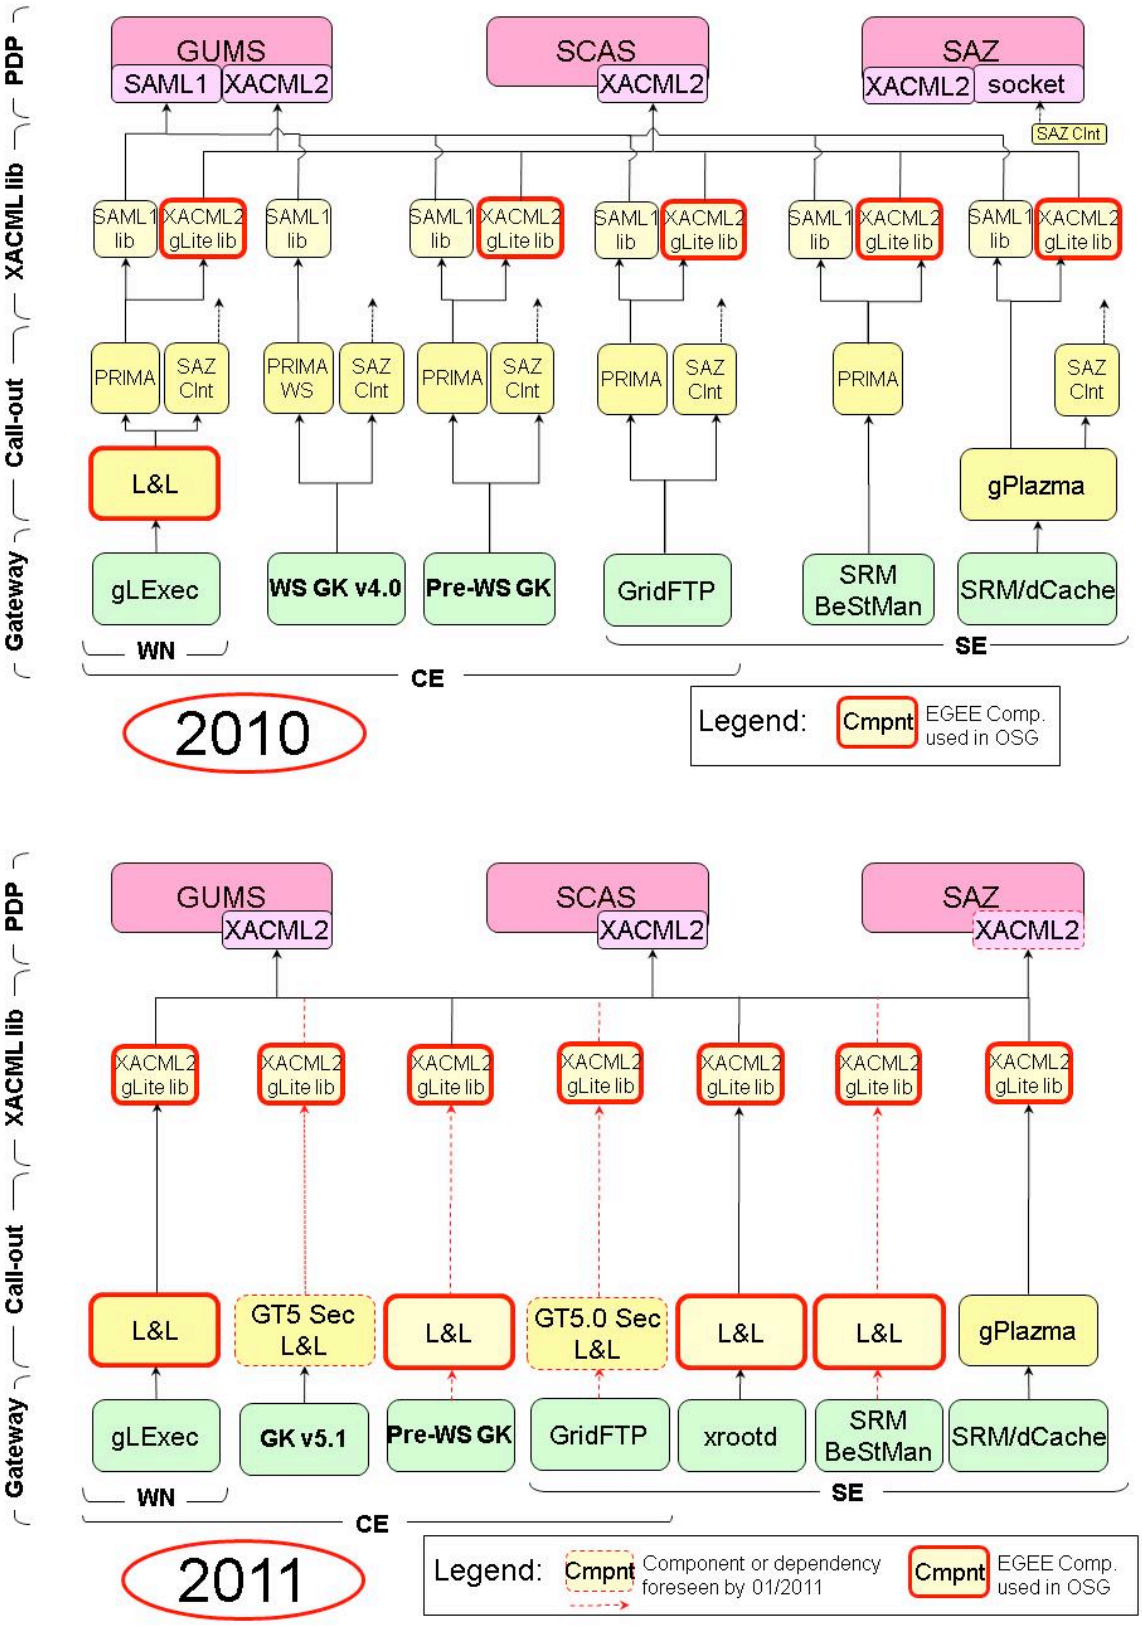

Figure 2. A diagram of the authorization call-out infrastructure for services commonly deployed in the Open Science Grid. Using common authorization modules drastically simplifies the infrastructure.

modules. This results in a complex infrastructure, to package, deploy, configure, maintain, and support. As the deployment of the authorization interoperability modules picks up (bottom diagram), the use of the LCAS/LCMAPS (L&L) framework as a common implementation simplifies the overall architecture.

In 2010 and 2011, the project focuses on integrating more resource-gateways with the Authorization Interoperability profile and on deploying the infrastructure beyond test beds (sec 4).

4. Deployment

The deployment of the L&L module in OSG is the key to reduce the complexity and maintenance load of the authorization infrastructure. The migration to the new infrastructure, however, requires packaging, testing, and administrative effort to simplify the software without offering any new functionality. Because of this, the deployment work has only recently started to get traction.

Pilot deployment sites are the University of Nebraska at Lincoln (UNL) and the Fermilab Campus Grid, FermiGrid [18]. In particular, the former is now enabling access to an Hadoop file system for all the Storage Element interfaces via XACML (SRM/BeStMan, GridFTP, xrootd). The latter is stress testing the XACML PDPs to enable an initial convergence of the GUMS and SAZ PEP modules to a common XACML implementation, PRIMA, in preparation for the future migration to L&L.

Making the deployment of the new modules easy is an important step to promote the adoption of the new infrastructure. We envision collaborating on the packaging of the authorization modules with the Virtual Data Toolkit [19], the *de facto* standard Grid middleware software stack.

5. Conclusions

EGEE and OSG have collaborated with the Globus Toolkit and Condor teams to release an Authorization Interoperability profile and XACML implementation document in 2008. The document describes a common "vocabulary" for OSG and EGEE to express the properties of authorization assertions. Authorization call-out module implementations are integrated with major resource gateways. The major advantages of the infrastructure are:

1. the ability to share and reuse software developed for EGI and OSG;
2. giving software providers reference protocols to integrate with both Grids infrastructures;
3. enabling the deployment of software developed by OSG or EGI in the EGI or OSG security infrastructures respectively. Because OSG and EGI are currently using different versions of the profile, "cross-deployment" cannot be achieved at this time anymore.

As of 2010, production deployments are slowly getting traction, despite the fact that the new software reduces the complexity of the site infrastructure without offering any additional functionality.

6. Acknowledgments

Fermilab is operated by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the United States Department of Energy. We thank Brian Bockelman at UNL for his work on the deployment of the Authorization Interoperability infrastructure.

7. References

[1] Pordes R, Petravick D, Kramer B, Olson D, Livny M, Roy A, Avery P, Blackburn K, Wenaus T, Wurthwein F, Foster I, Gardner R, Wilde M, Blatecky A, McGee, J, and Quick R. 2007. The Open Science Grid *Journal of Physics: Conference Series*, 78 15

[2] Laure E, Hemmer F, Aimar A, Barroso M, Buncic P, Di Meglio A, Guy L, Kunszt P, Beco S, Pacini F, Prelz F, Sgaravatto M, Edlund A, Mulmo O, Groep D, Fisher SM, and Livny M. 2004. Middleware for the next generation Grid infrastructure *Proceedings of Computing in High Energy Physics and Nuclear Physics* 2004, Interlaken, Switzerland 826

[3] Garzoglio G, Alderman I, Altunay M, Ananthakrishnan R, Bester J, Chadwick K, Ciaschini V, Demchenko Y, Ferraro A, Forti A, et al. 2009. Definition and Implementation of a SAML-

XACML Profile for Authorization Interoperability across Grid Middleware in OSG and EGEE *Journal of Grid Computing* DOI: 10.1007/s10723-009-9117-4
[4] Foster I and Kesselman C. 1997. Globus: A Metacomputing Infrastructure Toolkit *International Journal of Supercomputer Applications*, 11(2) 115-128
[5] Thain D, Tannenbaum T, and Livny M. 2005. Distributed Computing in Practice: The Condor Experience *Concurrency and Computation: Practice and Experience* 17(2-4) 323-356
[6] Garzoglio G et al. 2008. An XACML Attribute and Obligation Profile for Authorization Interoperability in Grids *Fremilab White Paper* CD-doc-2952-v2
[7] Cantor S, Kemp J, Philpott R, Maler R. 2005. Assertions and Protocols for the OASIS Security Assertion Markup Language (SAML) v2.0 *OASIS SSTC*
[8] Moses T et al. 2005. Extensible access control markup language (xacml) version 2.0 *Oasis Standard*
[9] Groep D. 2008. gLExec, SCAS and the way forward *Proceedings of the EGEE08 Conference - the Middleware Security Group, Istanbul, Turkey*
[10] Röblitz T, Schintke F, Reinefeld A, Barring O, Lopez M B, Cancio G, Chapeland S, Chouikh K, Cons L, Poznanski P, et al. 2004. Autonomic Management of Large Clusters and Their Integration into the Grid *Journal of Grid Computing* 2(3): 247-260
[11] Lorch M, Kafura D, Fisk I, Keahey K, Carcassi G, Freeman T, Peremutov T, Rana AS. 2005. Authorization and account management in the Open Science Grid *Proceedings of the 6th IEEE/ACM International Workshop on Grid Computing*, 2005
[12] Chadwick K, Sharma N, Timm SC, Yocum DR. 2009. FermiGrid – Site AuthoriZation (SAZ) Service *Proceedings of Computing in High Energy Physics and Nuclear Physics 2009, Prague, Czech Republic*
[13] N. Wang, G. Garzoglio, B. Ananthan, S. Timm, T. Levshina. 2010. Toward SVOPME, a Scalable Virtual Organization Privileges Management Environment, *Proceedings of the international Symposium on Grid Computing 2010, Taipei, Taiwan*
[14] The Joint Security Policy Group (JSPG): http://proj-lcg-security.web.cern.ch/proj-lcg-security Accessed Nov 2010
[15] Levshina T. 2006. The Virtual Organization Management Registration Service Proceedings of Computing in High Energy Physics and Nuclear Physics 2006, Mumbai, India
[16] Ceccanti A, Ciaschini V, Dimou M, Garzoglio G, Levshina T, Traylen S, Venturi V. 2009. VOMS/VOMRS Utilization patterns and convergence plan Proceedings of Computing in High Energy Physics and Nuclear Physics 2009, Prague, Czech Republic
[17] M. Altunay et. al., 2008 An XACML Attribute and Obligation Profile for Authorization Interoperability in Grids, *FNAL Doc DB 2685-v1, Fermilab*
http://cddocdb. fnal.gov/cgibin/ShowDocument?docid=268
[18] D.R. Yocum, E. Berman, P. Canal, K. Chadwick, T. Hesselroth, G. Garzoglio, T. Levshina, V. Sergeev, I. Sfiligoi, N. Sharma, and S. Timm, 2007, FermiGrid, *Proceedings of TeraGrid '07, Madison, Wisconsin*
[19] The Virtual Data Toolkit (VDT). Accessed on Nov 2010: http://vdt.cs.wisc.edu/